

REMARKS

This Amendment is filed in response to the Final Office action dated November 25, 2007 with a Request for Continued Examination and the associated fee. All objections and rejections are respectfully traversed.

Claims 1-9 and 13-35 are in the case.

Rejections Under 35 U.S.C. §103(a)

At paragraph 11 of the Final Office Action, the Examiner rejected claims 1-3, 13-14, and 35 under 35 U.S.C. § 103(a) as being unpatentable over Craddock et al., U.S. Publication No. 2003/0061296, published March 27, 2003, (hereinafter “Craddock”) in view of Pandya et al., U.S. Publication No. 2004/0037319, published February 26, 2004, (hereinafter “Pandya”).

Applicant's claimed novel invention, as set forth in representative claim 1, comprises in part:

1. A method for initiating a peer-to-peer communication session, the method comprising :
 - initiating a boot process;
 - initializing a cluster connection manager early in the booting process;
 - initiating, initiating by the cluster connection manager, a first remote direct memory access (RDMA) read operation directed to a cluster partner before a storage operating system executing on the cluster partner is fully active, the RDMA read operation bypassing the operating system;*
 - performing, in response to a successful first RDMA read operation, a first RDMA write operation to the cluster partner;
 - performing, in response to a successful RDMA write operation, a second RDMA read operation directed to the cluster partner; and
 - performing, in response to a successful second RDMA read operation, a second RDMA write operation to the cluster partner earlier in the booting process.

Craddock discloses a technique for processing storage I/O in a system area network (SAN). Craddock uses I/O transactions which represent a unit of I/O work and typically contain multiple messages. These I/O transactions are read from a specific disk sector into a specific host memory location. Craddock provides a mechanism for initiating and completing one or more I/O transactions using memory semantic messages which are transmitted by means of a remote direct access (RDMA) operation. Specifically, a process running on a host first reserves a memory space for holding read data. The process then invokes a device driver associated with the storage device adapter, specifying that data from the storage device is read into read data memory space. At the close of the read transaction, the adapter generates a response and an associated write RDMA with an immediate work queue element that is interpreted, processed, and transmitted via RDMA transfer to a host where it is stored in a location which was reserved for the response message.

Pandya discloses issuing “[a] TCP/IP processor and data processing engines for use in the TCP/IP processor.” See Abstract. Specifically, “FIG. 10 illustrates the TCP/IP stack functionality that is implemented in the described IP processor system. These functions provide an interface to the upper layer protocol functions to carry the IP storage traffic as well as other applications that can benefit from direct OS TCP/IP bypass, RDMA or network sockets direct capabilities or combination thereof to utilize the high performance TCP/IP implementation of this processor (i.e., not the processor or the operating system of the cluster partner). Pandya allows applications to access a processor directly by running the applications on an initiator. Running the applications on an initiator allows the application to register a region of memory that is then made available to its peers directly without substantial host intervention through normal RDMA protocol data transfer.

Applicant respectfully urges that neither Craddock nor Pandya show Applicant’s claimed novel use of “***initiating, initiating by the cluster connection manager, a first remote direct memory access (RDMA) read operation directed to a cluster partner be-***

fore a storage operating system executing on the cluster partner is fully active, the RDMA read operation bypassing the operating system.”

Applicant claims a system and method for establishing (i.e., initiating) a peer connection using reliable RDMA primitives (i.e., any protocol that supports specific ports and network addresses that support RDMA) early in a booting process. In Applicant's system, the RDMA read operation bypasses the operating system before the operating system on the cluster partner is fully active. That is, Applicant's invention uses specific ports and network addresses so that it knows that the connection will be reliable based on the protocols known to support RDMA operations. That is, it does not rely on the Virtual Interface Layer to make the determination. Specifically, the method initiates a peer-to-peer communication session by first attempting a remote direct memory access read operation directed to a predefined hardware address and a predefined port number which is known to support remote direct access memory operations. If a peer-to-peer connection is successful, then a RDMA write operation is performed wherein it too is directed to the predefined hardware address and the predefined port number.

Craddock does not disclose accessing in any data from any location initiating a remote RDMA read of a cluster partner before the operating system executing on the cluster partner is fully active. Rather Craddock discloses attempting to read data and write data merely using ordinary RDMA techniques. Craddock, like Pandya, merely reserves a memory space for holding read data and then specifies that data from a storage device is to be read into the reserved data memory space. Craddock makes no mention of Applicant's initiating the RDMA read operation on a cluster partner *early in the booting process before the storage operation system executing on the cluster partner is fully active*.

Additionally, Pandya does not initiate a read request directed to a cluster partner before a storage operating system executing on the cluster partner is fully active. Rather Pandya discloses merely directly accessing a region of memory without any host intervention. There is no mention in Pandya of initiating the RDMA on a cluster partner be-

fore the operating system on the cluster partner is active. Pandya only mentions in passing that [0101] an application can benefit from direct OS TCP/IP bypass but does not state or imply that the OS is not fully active at the time of the bypass. Applicant specifically claims bypassing the cluster partner's operating system and initiating the RDMA read before the operating system on the cluster partner is fully active.

Accordingly, the Applicant respectfully urges that the combination of Craddock and Pandya is legally insufficient to make obvious the present claims under 35 U.S.C. §103(a) because of the absence of the Applicant's claimed novel *"initiating, initiating by the cluster connection manager, a first remote direct memory access (RDMA) read operation directed to a cluster partner before a storage operating system executing on the cluster partner is fully active, the RDMA read operation bypassing the operating system."*

At paragraph 5 of the Final Office Action, the Examiner rejected claims 15-19 under 35 U.S.C. § 103(a) as being unpatentable over Craddock in view of Prakash et al., U.S. Patent No. 6,434,626, issued on August 13, 2002 (hereinafter "Prakash") and further in view of Pandya.

Applicant's claimed novel invention, as set forth in representative claim 15, comprises in part:

15. A method comprising :

initiating a boot process;

initializing a cluster connection manager early in the boot process;

initiating, early in the booting process, a peer-to-peer communication session, by a cluster connection manager, before a storage operating system executing on the cluster partner is fully active which bypasses an operating system on a storage system by attempting a first remote direct memory access read operation directed to a predefined hardware address and a predefined port number, the predefined hardware address and the predefined port number previously known to support a RDMA operation; and

performing early in the booting process, in response to a successful initiating, a first remote direct memory access write operation directed to the predefined hardware address and the predefined port number.

As discussed above, Craddock discloses a method, computer program product, and distributed data processing system for processing storage I/O in a system area network (SAN). Craddock uses ...memory semantic messages which are ...read from a specific disk sector...” Therefore, the system must be fully initiated before any I/O transactions can be read from a specific disk sector.

Pandya discloses issuing “[a] TCP/IP processor and data processing engines for use in the TCP/IP processor.” See Abstract. Specifically, “FIG. 10 illustrates the TCP/IP stack functionality that is implemented in the described IP processor system. These functions provide an interface to the upper layer protocol functions to carry the IP storage traffic as well as other applications that can benefit from direct OS TCP/IP bypass, RDMA or network sockets direct capabilities or combination thereof to utilize the high performance TCP/IP implementation of this processor (i.e., not the processor or the operating system of the cluster partner). Pandya allows applications to access a processor directly by running the applications on an initiator. Running the applications on an initiator allows the application to register a region of memory that is then made available to its peers directly without substantial host intervention through normal RDMA protocol data transfer.

Prakash merely discloses initiating a peer-to-peer communication session.

Applicant respectfully urges that neither Craddock nor Prakash show Applicant’s claimed novel use of “*initiating, early in the booting process, a peer-to-peer communication session, by a cluster connection manager, before a storage operating system executing on the cluster partner is fully active which bypasses an operating system on a storage system by attempting a first remote direct memory access read operation directed to a predefined hardware address and a predefined port number, the predefined*

hardware address and the predefined port number previously known to support a RDMA operation.”

As noted above, Applicant claims a system and method for establishing (i.e., initiating) a peer connection using reliable RDMA primitives (i.e., any protocol that supports specific ports and network addresses that support RDMA) *before a storage operating system executing on the cluster partner is fully active*. Thus, allowing a quicker connection time and faster RDMA operations.

Craddock does not disclose accessing any data from any location *before a storage operating system executing on the cluster partner is fully active*. Craddock transfers messages via the verbs interface and a message and data interface layer service. Craddock states in paragraph 0050 that this system uses both a verbs interface layer and a message and data interface layer to operate. These layers in combination act like a virtual interface layer (i.e., the operating system at full capacity) because they define the destinations (i.e., addresses and ports) which must be used in order for the RDMA operation to function properly. Applicant’s invention accesses the cluster partner *before a storage operating system executing on the cluster partner is fully active* (i.e., where the virtual interface layer runs). A

Additionally, Prakash merely discloses initiating a peer-to-peer communication session. Thus, Prakash does not disclose accessing the cluster partner *before a storage operating system executing on the cluster partner is fully active*

As noted above, Pandya only mentions in passing that [0101] an application can benefit from direct OS TCP/IP bypass but does not state or imply that the OS is not fully active at the time of the bypass. Applicant specifically claims bypassing the cluster partner’s operating system and initiating the RDMA read before the operating system on the cluster partner is fully active.

Accordingly, Applicant respectfully urges that Craddock and Pandya either taken singly or in any combination are legally insufficient to render the presently claimed in-

vention obvious under 35 U.S.C. § 103(a) because of the absence in each of the cited patents of Applicant's claimed *"initiating, early in the booting process, a peer-to-peer communication session, by a cluster connection manager, before a storage operating system executing on the cluster partner is fully active which bypasses an operating system on a storage system by attempting a first remote direct memory access read operation directed to a predefined hardware address and a predefined port number, the predefined hardware address and the predefined port number previously known to support a RDMA operation."*

At paragraph 17 of the Final Office Action, claims 4-9 were rejected under 35 U.S.C. § 103(a) over Craddock in view of Pandya even further in view of Costello et al., U.S. Patent Application Publication No. 2003/0078946 (hereinafter "Costello"). Claims 4-9 are believed to be dependent from allowable independent claim 1, and therefore in condition for allowance.

At paragraph 24 of the Office Action, the Examiner rejected claims 20-31 under 35 U.S.C. § 103(a) as being unpatentable over Prakash in view of Sutherland et al., U.S. Publication No. 2002/0114341, published on August 22, 2002 (hereinafter "Sutherland") and in further view of Pandya.

Applicant's claimed novel invention, as set forth in representative claim 20, comprises in part:

20. A system configured to establish reliable peer-to-peer communication among storage systems of a clustered environment, the system comprising:

a booting process executed by a processor;

a peer process executing on each storage system partner having an operating system; and

a cluster connection manager executing on each storage system partner, the cluster connection manager establishing a reliable peer-to-peer connection between each peer process early in the booting process before a storage operating system executing on a cluster partner is fully active by connecting to a predetermined port number using a predetermined network address, the reliable peer-to-peer connection bypassing the operating system.

As noted above, Pandya only mentions in passing that [0101] an application can benefit from direct OS TCP/IP bypass but does not state or imply that the OS is not fully active at the time of the bypass. Applicant specifically claims bypassing the cluster partner's operating system and initiating the RDMA read before the operating system on the cluster partner is fully active. Furthermore, Prakash merely initiates a peer-to-peer communication session.

Sutherland merely discloses a cluster connection manager.

Applicant respectfully urges that neither Craddock, Prakash, nor Sutherland either taken singly or in any combination are legally insufficient to render the presently claimed invention obvious under 35 U.S.C. § 103(a) because of the absence in each of the cited patents of Applicant's claimed use of ***"a cluster connection manager executing on each storage system partner, the cluster connection manager establishing a reliable peer-to-peer connection between each peer process early in the booting process before a storage operating system executing on a cluster partner is fully active by connecting to a predetermined port number using a predetermined network address, the reliable peer-to-peer connection bypassing the operating system."***

At paragraph 30 of the Final Office Action, claims 32 was rejected under 35 U.S.C. § 103(a) over Craddock, Prakash and Sutherland in view of "what was well known in the art." Claim 32 is believed to be dependent from allowable independent claim 28, and therefore in condition for allowance.

At paragraph 32 of the Final Office Action, the Examiner rejected claim 33 under 35 U.S.C. § 103(a) as being unpatentable over Craddock and Pandya in view of Boyd et al., U.S. Publication No. 2004/0049600, published on March 11, 2004 (hereinafter "Boyd"). At paragraph 34, of the Final Office Action, the Examiner rejected claim 34 under 35 U.S.C. § 103(a) as being unpatentable over Craddock, Pandya and Boyd in view of Sutherland.

Applicant's claimed novel invention, as set forth in representative claim 33, com-

prises in part:

33. A method comprising:

initiating a boot process;

initializing a cluster connection manager early in the boot process;

initializing, early in the booting process, a first remote direct memory access (RDMA) read operation that bypasses the operating system and is directed to a specific cluster partner before a storage operating system executing on a cluster partner is fully active d, using a specific port number and a specific address that support a RDMA operations; and

performing a second RDMA read operation directed to a specific cluster partner before a higher virtual interface layer has fully initialized, using a specific port number and a specific address that support a RDMA operations.

As discussed above, Craddock discloses a method, computer program product, and distributed data processing system for processing storage I/O in a system area network (SAN). Craddock uses ...memory semantic messages which are...read from a specific disk sector... Therefore, the system must be fully initiated before any I/O transactions can be read from a specific disk sector.

Pandya only mentions in passing that [0101] an application can benefit from direct OS TCP/IP bypass but does not state nor imply that the OS is not fully active at the time of the OS TCP/IP bypass.

Boyd discloses a memory management system which uses a virtual interface layer to assign virtual addresses to physical addresses via a memory table. The memory management system is connected to an IP net (Fig 1). The IP net allows the user to bypass the host processor node and directly access the RAID subsystem. Within the RAID subsystem is a memory and a processor, the memory contains an operating system which all accesses to the RAID system must still go through. That is, although the user can bypass the host processor's operating system, the user may still not bypass the operating system of the RAID subsystem.

Sutherland merely discloses utilizing a cluster connection manager.

Applicant respectfully urges that neither Craddock, Pandya, nor Boyd show Applicant's claimed novel use of *"initializing, early in the booting process, a first remote direct memory access (RDMA) read operation that bypasses the operating system and is directed to a specific cluster partner before a storage operating system executing on a cluster partner is fully active d, using a specific port number and a specific address that support a RDMA operations."*

As noted above, Applicant claims a system and method for establishing (i.e., initiating) a peer connection using reliable RDMA primitives (i.e., any protocol that supports specific ports and network addresses that support RDMA) *before a storage operating system executing on the cluster partner is fully active*. Thus, allowing a quicker connection time and faster RDMA operations.

Boyd merely discloses what a virtual interface layer does. Applicant is not claiming a virtual interface layer but rather a method for initiating an RDMA read operation on a cluster partner *before a storage operating system executing on a cluster partner is fully active*. Boyd allows the user to bypass the host processors operating system. However, as can be seen by Boyd's Figure 1, the user may not bypass the storage systems operating system *before a storage operating system executing on a cluster partner is fully active*. Thus, the user can only directly access an IPSOE (Figure 1, 172) while the operating system is fully functional.

Pandya does not rectify this deficiency. As noted above, Pandya only mentions in passing that [0101] an application can benefit from direct OS TCP/IP bypass but does not state or imply that the OS is not fully active at the time of the bypass. Applicant specifically claims bypassing the cluster partner's operating system and initiating the RDMA read before the operating system on the cluster partner is fully active. Furthermore, Prakash merely initiates a peer-to-peer communication session.

Applicant respectfully urges that neither Craddock, Prakash, nor Boyd either taken singly or in any combination are legally insufficient to render the presently claimed invention obvious under 35 U.S.C. § 103(a) because of the absence in each of the cited

patents of Applicant's claimed use *initializing, early in the booting process, a first remote direct memory access (RDMA) read operation that bypasses the operating system and is directed to a specific cluster partner before a storage operating system executing on a cluster partner is fully active d, using a specific port number and a specific address that support a RDMA operations.*

Conclusion

All independent claims are believed to be in condition for allowance.

All dependent claims are believed to be dependent from allowable independent claims, and therefore in condition for allowance.

Favorable action is respectfully solicited.

Please charge any additional fee occasioned by this paper to our Deposit Account No. 03-1237.

Respectfully submitted,

/Stephen D. LeBarron/
Stephen D. LeBarron
Reg. No. 62,479
CESARI AND MCKENNA, LLP
88 Black Falcon Avenue
Boston, MA 02210-2414
Telephone: (617) 951-2500
Facsimile: (617) 951-3927